Low power Architectures

Lecture #1:Introduction

Dr. Avi Mendelson

mendlson@ee.technion.ac.il
Contributors: Ronny Ronen, Eli Savransky, Shekhar Borkar, Fred Pollack

Technion, EE department

Dr. Avi Mendelson, Intel

Why Power meter?

Moore's Law

For many years technology obeys Moore's Law

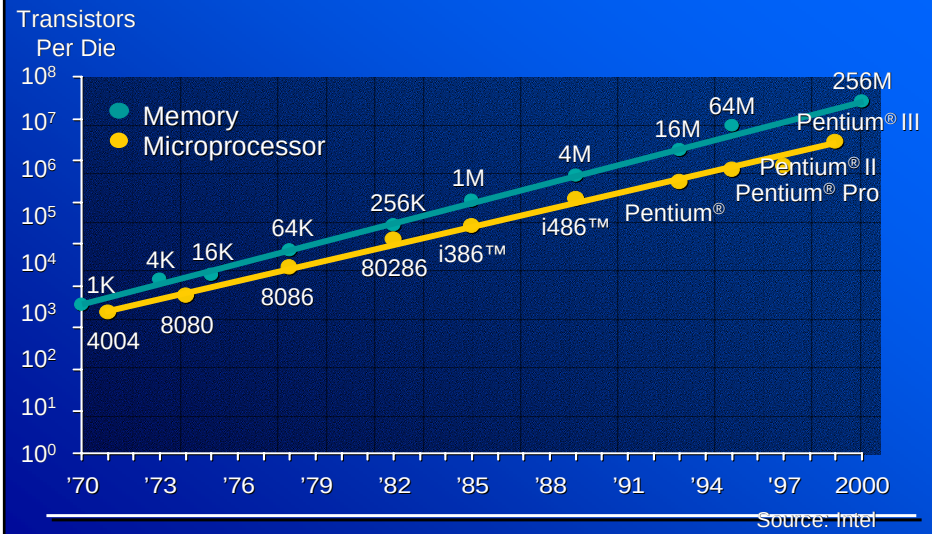

Source: Intel

In the Last 25 Years Life was Easy(*)

- Doubling of transistor density every 30 months
- Increasing die sizes, allowed by
  - Increasing Wafer Size
  - Process technology moving from "black art" to "manufacturing science"
- ⇒ *Doubling of transistors every 18 months*

Tech	Old µArch	mm (linear)	New µArch	mm (linear)	Ratio
1.0µ	i386C	6.5	i486	11.5	3.1
0.7µ	i486C	9.5	Pentium®	17	3.2
0.5µ	Pentium®	12.2	Pentium® Pro	17.3	2.1
0.18µ	Pentium® III	10.3	Next Gen	?	2--3

Implications: (in the same technology)

1. New µArch ~ 2-3X die area of the last µArch

2. Provides 1.5-1.7X integer performance of the last µArch

(*) *source* Fred Pollack, Micro-32

Suddenly, the power monster appears

The power crisis – power consumption

Environmental burden of CPUs!

- Total power consumption of CPUs in world's PCs:
  1992: 160 MWatts (87M CPUs)
  2001: 9,000 MWatts (500M CPUs)
- That's 4 Hoover Dams!

Courtesy: United States Department of the Interior
Bureau of Reclamation - Lower Colorado Region

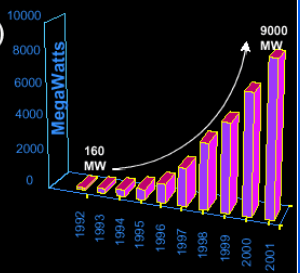

[Source: Dataquest (for installed base) + estimates for avg. installed CPU power]
Projected with PentiumII™ Power

Sourse: cool-chips, Micro 32

intel

Andy's vision: 1 Billion Connected PCs!

Processor Power Evolution

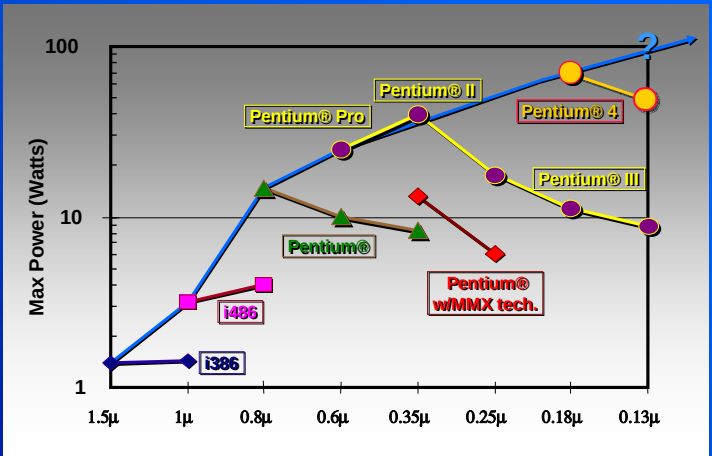

- *New generation: always increase power*
- *Compactions: higher performance at lower power*
- *One size fits all: start with high power segment and shrink it to Mobile*

The power crisis: *Power Density*
A real thread to the Moor law

- Think of watts/cm^2
- Power is not distributed evenly over the chip. A failure can happen if a single point reach the max power point.
- Complex algorithms lead to denser power:
  - Dense random logic
- Timing pressure leads to faster/bigger/power-hungrier gates
  - Designers put together units that communicate with each other. It creates “regions” with high activity factors -> hot spots.

Power Density

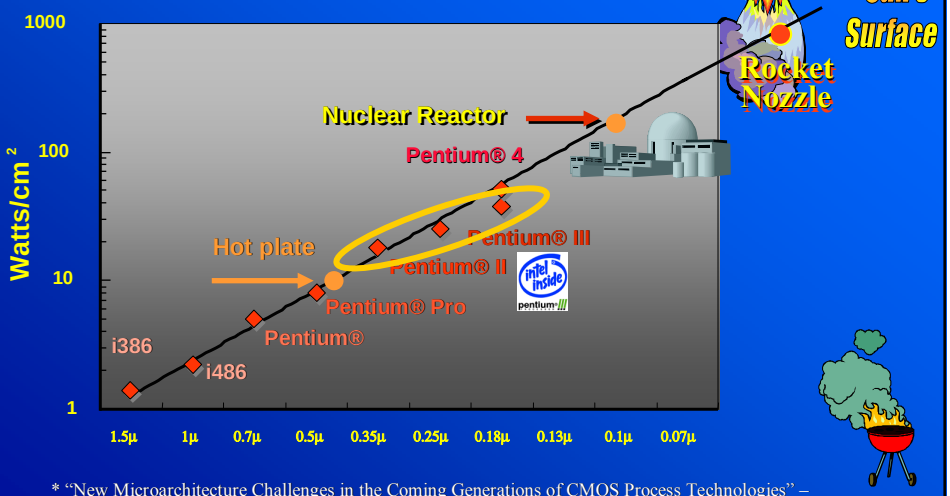

* "New Microarchitecture Challenges in the Coming Generations of CMOS Process Technologies" – Fred Pollack, Intel Corp. Micro32 conference key note - 1999.

Some implications

- We can't build microprocessors with ever increasing power density and die sizes
- The constraint is power – not manufacturability
- The design of any future micro-processor should take power into consideration. We need to distinguish between different aspects of power:
  - Power delivery
  - Max power (TJ)
  - Power density - hot spots
  - Energy – static + dynamic
- Power and Energy aware design should take care of each of these aspects

One size does not fit all anymore

Why power and power density increase over time

Basic terminology

Power and the digital world…

- **Power is consumed:**
  - When capacitance is charged and discharged
  - A charged cap is a logical '1', a discharged cap is '0'

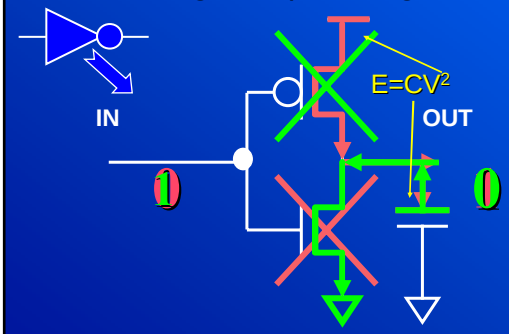

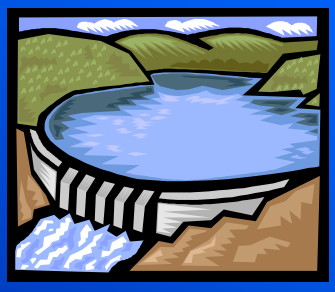

- The capacitance can be the gates of other transistors or wires (buses and long interconnects)

Power and the digital world (2)…

- Secondary effects like leakage and short-circuit current are increasing with advanced process technologies

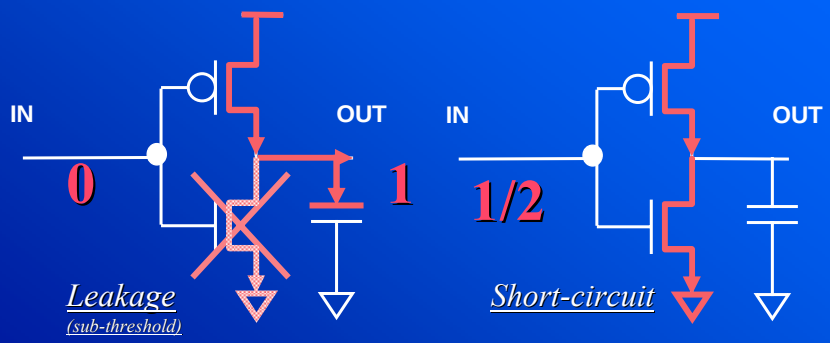

- Leakage is growing dramatically
  - 7% now, expect 20% in next process technology, 50% in next one
  - … Unless we do something (and we will)

Power & Energy

Energy

- "The capacity for doing work" *
- Important for
  - Battery life - lower energy per task ➔ longer battery life
  - Electric bills - lower energy per task ➔ lower bills
- Measured over time
- Proportional to the overall capacitance and to the voltage squared (CV^2)

* Merriam-Webster's Collegiate® Dictionary - http://www.m-w.com/

Power & Energy

Power

- *Work* done per time unit
  - Measured in Watts
- $P = \alpha CV^2f$
  (α: activity, C: capacitance, V: voltage, f: frequency)
- "Measured" at peak time
- Higher power ➔ higher current
  - Cannot exceed platform power delivery constrains
- Higher power ➔ higher temperature
  - Cannot exceed the thermal constrains

Voltage, Power, Frequency

- Transistor switches faster at higher voltage
  ➔ Higher voltage enables higher frequency
- Maximum frequency grows about linearly with voltage
  …Within a given voltage range Vmin-Vmax
  - V < Vmin
    ➔ transistors won't switch
  - V > Vmax
    ➔ the device may burn
- "The cube law":
  $P = kV^3$
  (or ~1%V=3%P)

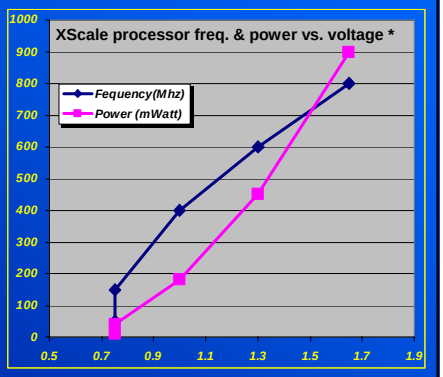

* Source: Intel Corp. (http://developer.intel.com)

The scalability theory – process technology

Target (Ideally):Eache generation (2-3 years)

- Reduce gate delay by 30% 50% freq gain
- 2. Increase density by 2x
  - 0.7 shrink on a side, 50% area reduction on compaction
  - Transistor Z and L shrink by 30%
  - Interconnect pitches shrink by ~30%
  - Add metal layers to make-up for (1) pitches < 30%, and (2) RC

Scaling theory--1 of 2

$Width = W = 0.7, Length = L = 0.7, t_{ox} = 0.7$

Lateral and vertical dimensions reduce 30%

$$Area\ \ Cap\ = C_a = \frac{0.7 \times 0.7}{0.7} = 0.7,$$

$$Fringing\ \ Cap\ = C_f = 0.7,$$

$$Total\ \ Cap\ \Rightarrow C = 0.7$$

Capacitance--area and fringing--reduce 30%

$$Die\ Area = X \times Y = 0.7 \times 0.7 = 0.7^2$$

Die area reduces 50%

Scaling theory--2 of 2

$$\frac{Cap}{Transistor} = \frac{0.7}{1} = 0.7$$

Capacitance per transistor reduces 30%

$$\frac{Cap}{Area} = \frac{0.7}{0.7 \times 0.7} = \frac{1}{0.7}$$

Capacitance per unit area increases 43%

$$Vdd = 0.7, V_t = 0.7, I = \frac{W}{t_{ox}}(Vdd - V_t) = \frac{0.7 \times 0.7}{0.7} = 0.7$$

$$T = \frac{C \times Vdd}{I} = \frac{0.7 \times 0.7}{0.7} = 0.7, Power = C \times V^2 \times f = \frac{0.7 \times 0.7^2}{0.7} = 0.7^2$$

Delay reduces 30%, power reduces 50%

Process Technology – the Enabler

- Every process generation (every 2-3 years), **Ideally:**
  - **Shorten gate delay by 30%**
    ➔ ~50% (100/70) frequency gain
  - **Vdd scaled down by ~30%**

➔ **Results:**

» $^2/_3$ reduction in energy/transition
(CV^2 ➔ $0.7 \times 0.7^2 = 0.34X$)

» $^1/_2$ reduction in power
(CV^2f ➔ $0.7 \times 0.7^2 \times 1.5 = 0.5X$)

» Power density unchanged

Ideal Scenarios...

- **Ideal "Shrink"**
  - *Same µarch*
  - 1X #Xistors
  - 0.5X size
  - 1.5X frequency
  - 0.5X power
  - 1X IPC (instr./cycle)
  - 1.5X performance
  - 1X power density
- **Ideal New µarch**
  - *Same die size*
  - 2X #Xistors
  - 1X size
  - 1.5X frequency
  - 1X power
  - 2X IPC
  - 3X performance
  - 1X power density

Looks good. Isn't it?

Process Technologies – Reality

- **But in reality:**
  - New designs squeeze frequency to 2X per process
  - New designs use more transistors (2X-3X to get 1.5X-1.7X perf)
- **So, every new process and architecture generation:**
  - Power goes up about 2X
  - Power density goes up 30%~80%
- **This is bad, and…**
- **Will get worse in future process generations:**
  - Voltage (Vdd) will scale down less
  - Leakage is going to the roof

Not as good as it first looked… Aha?

In Picture…

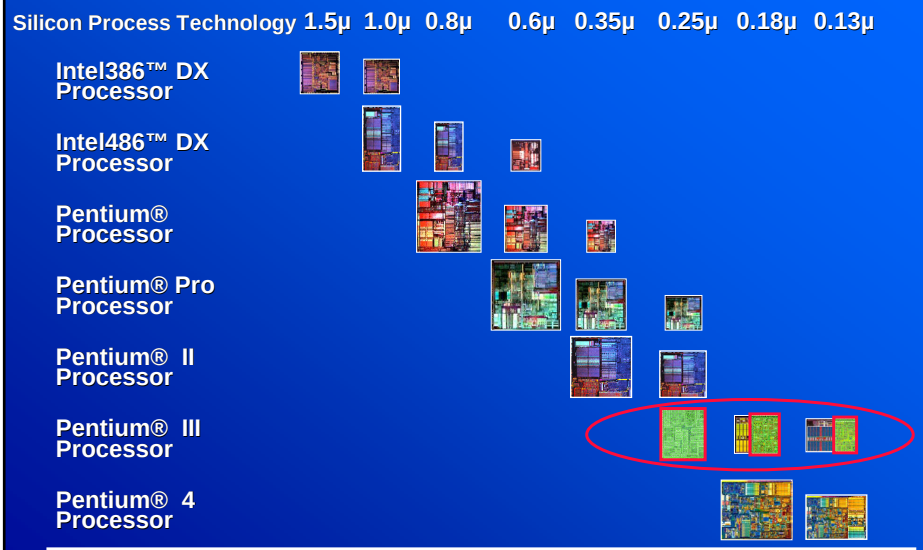

Performance Efficiency of µarchitectures

Tech	Old µArch	mm (linear)	New µArch	mm (linear)	Area Ratio	Perf Ratio
1.0µ	i386C	6.5	i486	11.5	**3.1**	
0.7µ	i486C	9.5	Pentium® proc	17	**3.2**	**1.8**
0.5µ	Pentium® proc	12.2	Pentium Pro® proc	17.3	**2.1**	**1.5**
0.18µ	Pentium III® proc	10.3	Pentium® 4 proc	14.7	**2**	**1.4**

Implications: (in the same technology)

1. New µarch ~2-3X die area of the last µarch
2. Provides 1.4-1.8X integer performance of the last µarch

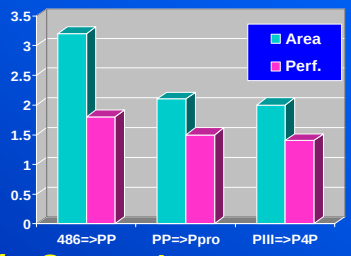

We are on the Wrong Side of a Square Law

Power Evolution *(Theoretical)*

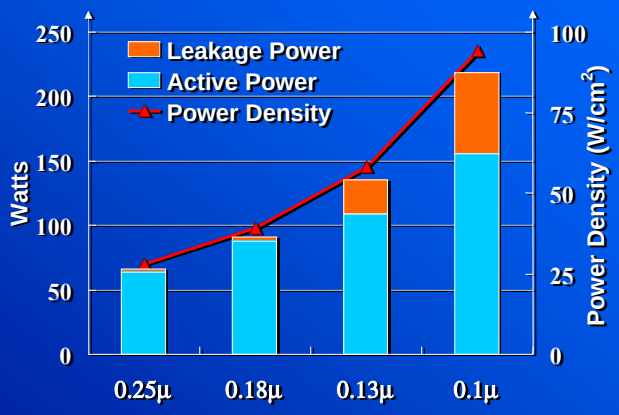

- For a 15mm/side die (225mm^2)
- Assume 2X frequency increase each generation
- Future process numbers are estimated

More aspects of power aware architectures

The "Power Bottleneck" (Hot Spots) The Thermal Story

Power Map

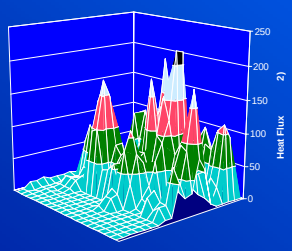

On-Die Temperature

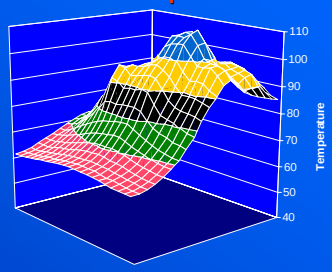

- Silicon is not a good heat conductor
- With high power density, cannot assume power uniformity
  - High temperature → high leakage → high power → higher temperature
- Artificially expanding the die size not help. Must attack the hot spots
- Smart layout that separates the hot units increases the processor's power envelop!

* "New Microarchitecture Challenges in the Coming Generations of CMOS Process Technologies" – Fred Pollack, Intel Corp. Micro32 conference key note - 1999.

The Mythical Power Envelop

- **CPU "power envelop"**
  - Maximum power that commercial cooling technologies can dissipate
- **Limited by**
  - Total system power
  - Processor power. Typical figures:
    Server <130W, Desktop 50-80W, Notebook 20-30W, sub-notebook <10W
- **Bigger systems cool better and dissipate more power**
  - Heat syncs
  - Heat pipes
  - Better TIM (Thermal Interface Materials)
- **Average power density matters:**
  - Uniformly distributed power allows for higher CPU dissipation

Energy Efficiency

- **Energy per task**
  - Proportional to # of processed insts. per task
  - Proportional to the average work consumed per instruction
  - Deteriorates as speculation increases and complexity grows

  Or Formally, per a given task,
  - Energy per retired instruction is: **β*W**, where
    - » β: Ratio of Total to Retired number of processed instructions
    - » W: Average energy spent in processing an instruction

    Both figures grow with every new micro-architecture
- ***In that respect:***
  high performance modern micro-architectures are less energy-efficient
- **Luckily, process technology offsets that by reducing energy per switch**

Voltage Scaling

- **Within a given voltage range, higher voltage allows higher freq.**
- **Used for trading power and frequency. Either**
  - Statically, at manufacturing time
  - Dynamically, at run time (e.g., Intel's SpeedStep® Technology
- **Actual range depends on specific design and process technology Examples*:**
  - Intel® XScale™ processors runs from 0.75V (150MHz/50mW) to 1.65V (800MHz/900mW)
  - Intel mobile Pentium® III processor sells from 1.1V (600MHz) to 1.7V (1GHz)

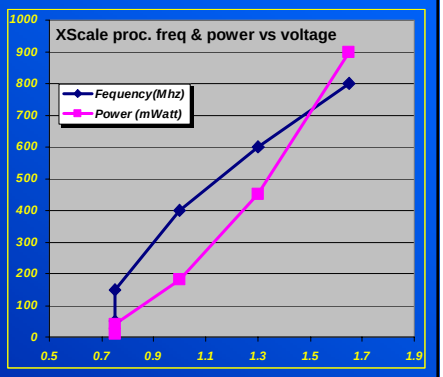

* Source: Intel Corp. (http://developer.intel.com)

Alpha hot spots

21064 Thermal Plot

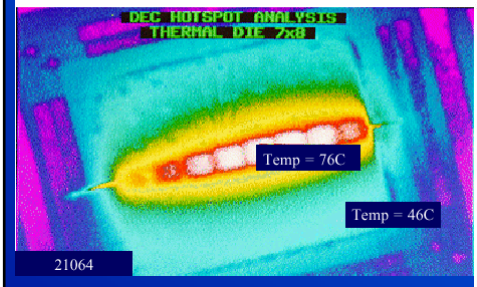

21164 Thermal Plot

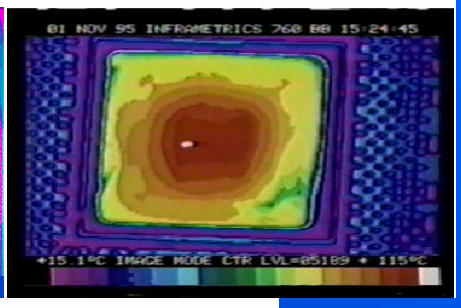

	Power (Watts)	Freq. (MHz.)	Die Size (mm^2)	Vdd
Alpha 21064	30	200	234	3.3
Alpha 21164	50	300	299	3.3

Area 30%

Freq. 50%

Power 67%

Source - CoolChips-99

Voltage Scaling (cont.)

- **Huge effect on Power:**
  - 20% freq reduction ➔ 20% voltage reduction
  - ➔ 35% energy reduction. ($\alpha CV^2 = \alpha C*0.8^2 = \alpha C*0.64$)
  - ➔ 50% power reduction. ($\alpha CV^2f = \alpha C*0.8^3 = \alpha C*0.51$)
- **Even more impressive if we recall:**
  - 20% freq hit ➔ only 10%-15% performance hit*
- ***Voltage scaling can be used to trade performance with power!***

* Depends mainly on core to bus frequency ratio and caches size.

What is the impact of the computer architecture?

Power and performance--trade-off

Voltage scaling alone is not enough to cap power

So far, the analysis indicates that:

You will have to tradeoff performance for power

- Reduce die area => Reduce Active C => Reduce Perf(C) & Power(C)
- Reduce Vdd and freq => Perf(freq) & Power(freq,Vdd^2)--cubic
- Reduce dies size, Vdd, and freq

Argument:

- **Tradeoff performance for power ➡ does Goal #1 make sense?**
- **Set a goal that comprehends both: performance and power.**
- **Use energy delay product (E*D) to evaluate tradeoffs**

E*D product (lower is better) may provide a better criteria

- E = energy / instruction
- = Power * sec / instruction
- = Watt / MIPS
- D = sec / instruction
- = 1 / MIPS
- E *D ~ Watt / $MIPS^2$

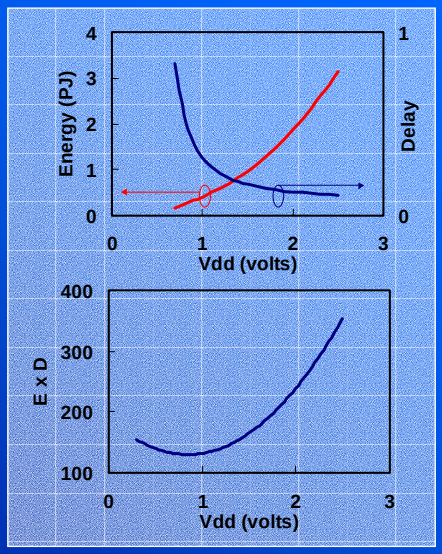

Power segments

One size does not fit all:

- Embedded systems
  - Most of the power is consumed by the CPU
  - We are not thermally limited.
  - What we really care about is battery life.
  - In real time systems we can take advantage of known "deadlines"
- Laptops (Mobile systems)
  - We are thermally limited.
  - We can not use deadlines (most of the time).
  - We need to optimize for max battery life and max performance in a given power envelop.
- Desktops:
  - We mainly care about power awareness and Thermal issues

Example: Adder Designs

- Various algorithms exist to implement an integer adder
  - Ripple, select, skip (x2), Look-ahead, conditional-sum.
  - Each with its own characteristics of timing and power consumption.

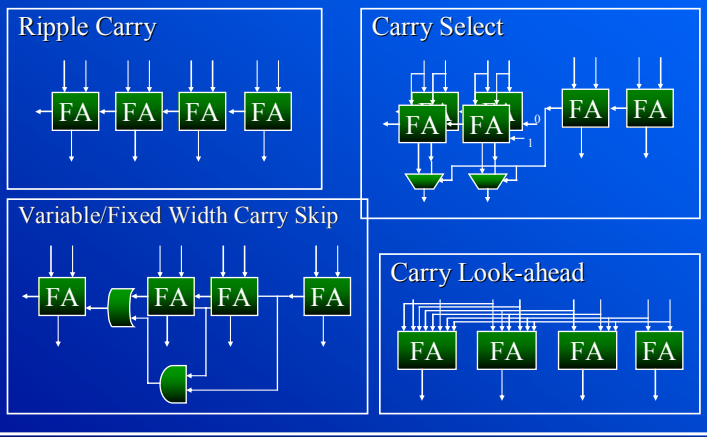

Power and Delay Numbers

- According to Callaway and Swartzlander*:

	Energy (pJ)	Delay (nSec)
Ripple Carry	117	54.27
Constant Width Carry Skip	109	28.38
Variable Width Carry Skip	126	21.84
Carry Lookahead	171	17.13
Carry Select	216	19.56
Conditional Sum	304	20.05

- If we must choose one option, as-is:
  - If power is the objective – use "constant width carry skip"
  - If delay is most important – use "carry look-ahead"

* "Estimating the power consumption of CMOS adders" - Callaway, T.K.; Swartzlander, E.E., Jr. 11th Symposium on Computer Arithmetic, 1993. Proceedings.

Power Complexity Metrics

- Power α C V^2 f
- Metrics: suppose we introduce new feature that consumes extra x power and gain y performance:
  1. Power/Perf (➔ Energy), assuming same technology (same C) and same voltage
     - » For battery life, energy bills.
     - » For a given power envelope – without voltage scaling.
  2. Power/$Perf^2$ (➔ Energy*Delay)
     - » Balance performance and power needs.
  3. Power/$Perf^3$ (➔ Energy*$Delay^2$)
     - » For a given power envelope – with voltage scaling. assuming that we can (1) trade frequency and voltage scaling, and (2) we can lower the voltage as much as we wish